



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
-----------------	-------------	----------------------	---------------------	------------------

10/644,133

08/20/2003

Chris P. Karamatas

BEA920030013US1

5334

61780 7590 07/09/2008
LAW OFFICES OF MICHAEL DRYJA
1474 N COOPER RD #105-248
GILBERT, AZ 85233

EXAMINER

SPITTLE, MATTHEW D

ART UNIT

PAPER NUMBER

2111

NOTIFICATION DATE

DELIVERY MODE

07/09/2008

ELECTRONIC

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Notice of the Office communication was sent electronically on above-indicated "Notification Date" to the following e-mail address(es):

MIKE@DRYJAPAT.COM

DETAILED ACTION

Claims 1, 4 – 10, 12 – 19, 21 – 25, and 28 – 30 have been examined.

Claim Rejections - 35 USC § 112

5 The following is a quotation of the second paragraph of 35 U.S.C. 112:

The specification shall conclude with one or more claims particularly pointing out and distinctly claiming the subject matter which the applicant regards as his invention.

 Claims 1, 7, 16, 21 and 25 are rejected under 35 U.S.C. 112, second paragraph,
10 as being indefinite for failing to particularly point out and distinctly claim the subject
matter which applicant regards as the invention.

 Claims 1, 7, 16, 21 and 25 recite, "...even if the second node has a cache,
memory, and at least one processor. A single claim which claims both an apparatus
and the method steps of using the apparatus is indefinite under 35 U.S.C. 112, second
15 paragraph. *> IPXL Holdings v. Amazon.com, Inc., 430 F.2d 1377, 1384, 77 USPQ2d
1140, 1145 (Fed. Cir. 2005);<Ex parte Lyell, 17 USPQ2d 1548 (Bd. Pat. App. & Inter.
1990).

Claim Rejections - 35 USC § 101

20 35 U.S.C. 101 reads as follows:

Whoever invents or discovers any new and useful process, machine, manufacture, or composition of matter, or any new and useful improvement thereof, may obtain a patent therefor, subject to the conditions and requirements of this title.

25 The claimed invention is directed to non-statutory subject matter. The
claim is directed to neither a "process" nor a "machine," but rather embraces or overlaps

two different statutory classes of invention set forth in 35 U.S.C. 101 which is drafted so as to set forth the statutory classes of invention in the alternative only

30 ***Claim Objections***

Claims 1, 7, 16, 21 and 25 are objected to under 37 CFR 1.75(c), as being of improper dependent form for failing to further limit the subject matter of a previous claim. Applicant is required to cancel the claim(s), or amend the claim(s) to place the claim(s) in proper dependent form, or rewrite the claim(s) in independent form. Claims 35 1, 7, 16, 21 and 25 recite the limitation, "...even if the second (third, fourth, etc) node has a cache, memory, and at least one processor..." This limitation is essentially the same as "whether or not" and thus fails to further limit the claim, since the step is carried out regardless.

40 ***Claim Rejections - 35 USC § 103***

The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

45 (a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

50 The factual inquiries set forth in *Graham v. John Deere Co.*, 383 U.S. 1, 148 USPQ 459 (1966), that are applied for establishing a background for determining obviousness under 35 U.S.C. 103(a) are summarized as follows:

1. Determining the scope and contents of the prior art.

- 55 2. Ascertaining the differences between the prior art and the claims at issue.
 3. Resolving the level of ordinary skill in the pertinent art.
 4. Considering objective evidence present in the application indicating
 obviousness or nonobviousness.

 Claims 1, 4 – 10, 12 – 19, and 21 – 30 are rejected under 35 U.S.C. 103(a) as
60 being unpatentable over Kiick (U.S. Pub. 2003/0200250) and what was well known in
 the art as evidenced by Rowlands et al. (U.S. 6,965,973), Fischer et al. (U.S.
 6,438,672), Drott et al. (U.S. 6,170,025), Agatsuma et al. (U.S. 7,237,099), and Chi et
 al. (U.S. 6,209,086).

 Regarding claim 1, Kiick teaches a method comprising at least one of:

65 Assigning interrupts for a plurality of input/output (I/O) devices among a plurality
 of nodes of a system based on at least one of: the nodes to which the I/O devices are
 connected; the nodes at which interrupt service routines for the I/O devices reside; and
 processors of the nodes for the nodes having processors, where one or more of the
 nodes have processors and memory (Paragraph 34 describes that interrupts should be
70 assigned to the “closest” processors, and not across node boundaries. Examiner
 interprets this to mean the interrupts for the I/O devices should be assigned to nodes to
 which they are connected or to nodes where the ISRs for the said I/O devices reside.);

 For each node of the system having processors, assigning the interrupts for the
 devices that are performance critical and that have been assigned to the node among
75 the processors of the node in a round-robin manner (Examiner interprets all devices in
 the reference to be considered “performance critical”; Paragraph 26);

Dynamically modifying assignments of the interrupts among the nodes of the system based on actual performance characteristics of the assignments (Paragraphs 25, 28, 31);

80 For each node of the system having processors, dynamically modifying assignments of the interrupts that are performance critical and that have been assigned to the node among the processors of the node based on actual performance characteristics of the assignments (Paragraphs 25, 28, 31).

Kiick fails to teach where one or more of the nodes are memoryless.

85 Examiner takes Official Notice that nodes in a NUMA system may or may not contain memory. Rowlands et al. evidences this (col. 4, lines 1 – 2).

Therefore, it would have been obvious to one of ordinary skill in this art at the time of invention by Applicant to provide memoryless nodes within the system of Kiick, since this is routine within NUMA systems.

90 Kiick fails to teach assigning the interrupts for the I/O devices to a first node having a cache, memory, and at least one processor and assigning the interrupt for the I/O device to the first node. Similarly, Kiick fails to teach repeating the assignment for the second and third nodes, respectively, if the first node to which the device is connected does not have a cache, memory, and at least one processor.

95 Examiner takes Official Notice that it would be obvious to assign the interrupt to a node which has a processor and memory, since these elements are required to service an interrupt (as evidenced by Agatsuma et al.; col. 4, line 64 - col. 5, line 9; col. 6, lines 15 - 22; Drott et al.; col. 21, lines 15 - 22; col. 22, lines 44 - 58; col. 23, lines 22 – 30;

and Fischer et al.; col. 9, lines 35 - 51). Examiner notes that all modern processor
100 architectures contain a "memory" since they contain registers which are operated upon
internally. Cache is old and notoriously well known in the art for providing improved
performance in computing systems (as evidenced by Chi et al.; col. 2, lines 29 – 36; and
Fischer et al.; col. 1, line 31 – col. 2, line 40).

Therefore, it would have been obvious to one of ordinary skill in this art at the
105 time of invention by Applicant to assign the interrupt for the I/O devices of Kiick to nodes
which contain processors and memory, since both of these components are known in
the prior art to be required in order to service an interrupt. It would have been obvious
to use choose nodes which have cache, since it is well-established in this art that
caching provides greater performance.

110
With regard to claim 4, Kiick teaches the method of claim 1, wherein dynamically
modifying the assignments of the interrupts among the nodes of the system comprises:

Measuring responsiveness of the node in processing the interrupt (paragraphs
27 – 30);

115 Kiick fails to explicitly teach assigning the interrupt to the node at which the
interrupt service routine for the I/O device resides; measuring responsiveness of the
node at which the interrupt service routine for the I/O device resides in processing the
interrupt; and where the responsiveness of the node to which the I/O device is
connected is better than the responsiveness of the node at which the interrupt service

120 routine for the I/O device resides, reassigning the interrupt to the node to which the I/O device is connected.

Kiick does, however, teach that the dynamic interrupt distributor should be aware of the system architecture, and re-assign interrupts to the “closest” processors (paragraph 34). Therefore, it would have been obvious to one of ordinary skill in this art
125 at the time of invention by Applicant to give preference in assigning interrupts to nodes at which the ISR for the device resides, or at which the I/O device itself resides, and then to re-assign interrupts accordingly to which configuration produced better performance as described in paragraphs 28 – 30).

130 With regard to claim 5, Kiick teaches the method of claim 4, wherein dynamically modifying the assignments of the interrupts among the nodes of the system comprises, for each assignment of an interrupt for an I/O device to a node, where the node is that at which the interrupt service routine for the I/O device resides:

Measuring responsiveness of the node in processing the interrupt (paragraphs
135 27 – 30);

Kiick fails to explicitly teach assigning the interrupt to the node to which the I/O device is connected; measuring responsiveness of the node to which the I/O device is connected in processing the interrupt; and where the responsiveness of the node at which the interrupt service routine for the I/O device is connected is better than the
140 responsiveness of the node to which the I/O device is connected, reassigning the interrupt to the node at which the interrupt service routine for the I/O device resides.

Kiick does, however, teach that the dynamic interrupt distributor should be aware of the system architecture, and re-assign interrupts to the “closest” processors (paragraph 34). Therefore, it would have been obvious to one of ordinary skill in this art at the time of invention by Applicant to give preference in assigning interrupts to nodes at which the I/O device is connected, or at which the ISR for the I/O device itself resides, and then to re-assign interrupts accordingly to which configuration produced better performance as described in paragraphs 28 – 30).

With regard to claim 6, Kiick teaches the method of claim 1, wherein for each node of the system, dynamically modifying the assignments of the interrupts that are performance critical and that have been assigned to the node among the processors of the node comprises:

Measuring the responsiveness of the processors of the node in processing the interrupts assigned thereto (paragraphs 27, 28, 35);

Where a differential between a best responsiveness and a worst responsiveness is greater than a threshold (paragraph 28; where a threshold may be interpreted as a “large enough difference”);

Reassigning at least one of the interrupts assigned to the processor having the worst responsiveness to the processor having the best responsiveness (paragraphs 27 – 30, 35).

With regard to claim 7, Kiick teaches a non-uniform memory access (NUMA) system comprising:

165 A plurality of nodes (Figure 1, items 102A, 102B);

A plurality of input/output (I/O) devices, each I/O device connected to one of the plurality of nodes and having an interrupt (Figure 1, items 110A, 110B);

An interrupt-assignor responsive to the I/O devices and the nodes to assign the interrupt for each I/O device to one of the plurality of nodes in a performance-optimized
170 manner (where an interrupt-assignor may be interpreted as a dynamic interrupt distributor; Figure 2, item 210; paragraphs 25, 28).

Kiick teaches a multiprocessor system that is tightly-coupled, and could have shared main memory, mass storage, and cache, and runs a single copy of an operating system (paragraph 7). These limitations define a NUMA system as evidenced by the
175 definition fromt5 Whatis.com, and therefore, Kiick implicitly describes a NUMA system for use with his invention.

Kiick fails to teach where one or more of the nodes are memoryless.

Examiner takes Official Notice that nodes in a NUMA system may or may not contain memory. Rowlands et al. evidences this (col. 4, lines 1 – 2).

180 Therefore, it would have been obvious to one of ordinary skill in this art at the time of invention by Applicant to provide memoryless nodes within the system of Kiick, since this is routine within NUMA systems.

Kiick fails to teach assigning the interrupts for the I/O devices to a first node having a cache, memory, and at least one processor and assigning the interrupt for the

185 I/O device to the first node. Similarly, Kiick fails to teach repeating the assignment for the second and third nodes, respectively, if the first node to which the device is connected does not have a cache, memory, and at least one processor.

Examiner takes Official Notice that it would be obvious to assign the interrupt to a node which has a processor and memory, since these elements are required to service
190 an interrupt (as evidenced by Agatsuma et al.; col. 4, line 64 - col. 5, line 9; col. 6, lines 15 - 22; Drott et al.; col. 21, lines 15 - 22; col. 22, lines 44 - 58; col. 23, lines 22 - 30; and Fischer et al.; col. 9, lines 35 - 51). Examiner notes that all modern processor architectures contain a "memory" since they contain registers which are operated upon internally. Cache is old and notoriously well known in the art for providing improved
195 performance in computing systems (as evidenced by Chi et al.; col. 2, lines 29 - 36; and Fischer et al.; col. 1, line 31 - col. 2, line 40).

Therefore, it would have been obvious to one of ordinary skill in this art at the time of invention by Applicant to assign the interrupt for the I/O devices of Kiick to nodes which contain processors and memory, since both of these components are known in
200 the prior art to be required in order to service an interrupt. It would have been obvious to use choose nodes which have cache, since it is well-established in this art that caching provides greater performance.

With regard to claim 8, Kiick teaches the system of claim 7, wherein the memory
205 of each node that has memory is local to the node and remote to all other nodes (Figure 1, items 108A, 108B; paragraph 23 describes each domain having domain-specific

Art Unit: 2111

memory (where a domain may be interpreted as a node, as described earlier in paragraph 23)), and the interrupt-assignor is to assign the interrupt for each I/O device to one of the plurality of nodes that has memory and at least one processor (where an
210 interrupt-assignor may be interpreted as a dynamic interrupt distributor; Figure 2, item 210; paragraphs 25, 28; all nodes (items 102A, 102B are shown in Figure 1 to have memory and at least one processor).

With regard to claim 9, Kiick teaches the system of claim 8, wherein at least one
215 of the I/O devices are performance critical, the interrupt-assignor further to assign the interrupt for each I/O device that is performance critical among the at least one processor of the node to which the interrupt has been assigned in a round-robin manner (Examiner interprets all of the I/O devices of the invention of Kiick to be performance critical, thereby necessitating the use of his invention to improve performance;
220 Paragraphs 25, 26, 28).

With regard to claim 10, Kiick describes the system of claim 7, wherein, for each node that has processors, the interrupt-assignment software is further to dynamically modify assignments of the interrupts that are performance critical among the at least
225 one processor of the node based on actual performance characteristics of the assignments. Examiner believes Applicant meant to refer to “the interrupt-assignor” instead of “interrupt-assignment software.” (Paragraphs 26, 28, 31).

With regard to claim 12, Kiick describes the system of claim 7, wherein the
230 interrupt-assignor is further to dynamically modify assignments of the interrupts among
the plurality of nodes based on actual performance characteristics of the assignments
(Paragraphs 26, 28, 31).

Regarding claim 13, Kiick teaches wherein the interrupt-assignor (paragraph 34,
235 where a interrupt-assignor may be interpreted as a dynamic interrupt distributor) is to
give primary preference in assigning the interrupt for each I/O device to the node to
which the I/O device is connected (paragraph 34, where a domain may be interpreted
as a node) where the node to which the I/O device is connected has a cache
(paragraph 10; Examiner interprets the processors within the processor complex (106A,
240 106B) as having on-chip cache), memory (Figure 1, items 108A, 108B) , and at least
one processor (Figure 1, items 106A, 106B).

Regarding claim 14, Kiick teaches wherein each I/O device further has an
interrupt service routine residing at one of the plurality of nodes, and the interrupt-
245 assignor (paragraph 34, where a interrupt-assignor may be interpreted as a dynamic
interrupt distributor) is to give secondary preference in assigning the interrupt for each
I/O device to the node at which the interrupt service routine of the I/O device resides
(paragraphs 28, 34; Examiner notes that paragraph 28 identifies re-assigning interrupts
to be equivalent to re-assigning ISRs) where the node at which the interrupt service
250 routine of the I/O device resides has a cache (paragraph 10; Examiner interprets the

processors within the processor complex (106A, 106B) as having on-chip cache),
memory (Figure 1, items 108A, 108B), and at least one processor ((Figure 1, items
106A, 106B).

255 With regard to claim 15, Kiick describes the system of claim 7, wherein the
interrupt-assignor resides within one of the plurality of nodes (where an interrupt-
assignor may be interpreted as a dynamic interrupt distributor; Figure 2, item 210;
paragraph 28 describes a predetermined processor in a domain (node) dedicated to run
the interrupt-assignor).

260

With regard to claim 16, Kiick teaches a computer-readable storage medium:

A computer readable medium;

Means in the medium for assigning interrupts for a plurality of input/output (I/O)
devices (paragraph 28 describes a dynamic interrupt distributor embodied as a program
265 module. Examiner identifies that a program module must be embodied on a computer
readable medium in order to be useful, and therefore implicitly describes this limitation)
among a plurality of nodes based on at least one factor selected from the set consisting
of: the nodes to which the devices are connected, and the nodes at which interrupt
service routines for the I/O devices reside (Paragraph 34 describes that interrupts
270 should be assigned to the “closest” processors, and not across node boundaries.
Examiner interprets this to mean the interrupts for the I/O devices should be assigned to

nodes to which they are connected or to nodes where the ISRs for the said I/O devices reside.), where one or more of the nodes have processors and memory.

Kiick fails to teach where one or more of the nodes are memoryless.

275 Examiner takes Official Notice that nodes in a NUMA system may or may not contain memory. Rowlands et al. evidences this (col. 4, lines 1 – 2).

Therefore, it would have been obvious to one of ordinary skill in this art at the time of invention by Applicant to provide memoryless nodes within the system of Kiick, since this is routine within NUMA systems.

280 Kiick fails to teach assigning the interrupts for the I/O devices to a first node having a cache, memory, and at least one processor and assigning the interrupt for the I/O device to the first node. Similarly, Kiick fails to teach repeating the assignment for the second and third nodes, respectively, if the first node to which the device is connected does not have a cache, memory, and at least one processor.

285 Examiner takes Official Notice that it would be obvious to assign the interrupt to a node which has a processor and memory, since these elements are required to service an interrupt (as evidenced by Agatsuma et al.; col. 4, line 64 - col. 5, line 9; col. 6, lines 15 - 22; Drott et al.; col. 21, lines 15 - 22; col. 22, lines 44 - 58; col. 23, lines 22 – 30; and Fischer et al.; col. 9, lines 35 - 51). Examiner notes that all modern processor
290 architectures contain a "memory" since they contain registers which are operated upon internally. Cache is old and notoriously well known in the art for providing improved performance in computing systems (as evidenced by Chi et al.; col. 2, lines 29 – 36; and Fischer et al.; col. 1, line 31 – col. 2, line 40).

Therefore, it would have been obvious to one of ordinary skill in this art at the
295 time of invention by Applicant to assign the interrupt for the I/O devices of Kiick to nodes
which contain processors and memory, since both of these components are known in
the prior art to be required in order to service an interrupt. It would have been obvious
to use choose nodes which have cache, since it is well-established in this art that
caching provides greater performance.

300
With regard to claim 17, Kiick teaches the medium of claim 16, wherein the
means is for assigning the interrupts among the plurality of nodes further based on
whether the nodes have processors and memories (Kiick describes assigning ISRs to
processors which have associated memories; paragraph 14).

305
With regard to claim 18, Kiick describes the medium of claim 16, wherein the
means, for each node having processors, is further for assigning the interrupts for the
devices that are performance critical and that have been assigned to the node among
the processors of the node in a round-robin manner (Examiner interprets all of the I/O
310 devices of the invention of Kiick to be performance critical, thereby necessitating the
use of his invention to improve performance; Paragraphs 25, 26, 28).

With regard to claim 19, Kiick describes the medium of claim 18, wherein the
means, is further for dynamically modifying assignments of the interrupts among the
315 nodes based on actual performance characteristics of the assignments, and, for each

node having processors, for dynamically modifying assignments of the interrupts that are performance critical and that have been assigned to the node among the processors of the node based on actual performance characteristics of the assignments (paragraphs 25, 28, 31).

320

With regard to claim 21, teaches describes a computer-readable storage medium comprising:

An interrupt-assignor (Figure 2, item 210; paragraph 28) to assign interrupts for a plurality of input/output (I/O) devices among a plurality of nodes based on at least one
325 factor selected from the set consisting of:

The nodes to which the I/O devices are connected;

The nodes at which interrupt service routines for the I/O devices reside, where one or more of the nodes have processors and memory.

(Paragraph 34 describes that interrupts should be assigned to the “closest”
330 processors, and not across node boundaries. Examiner interprets this to mean the interrupts for the I/O devices should be assigned to nodes to which they are connected or to nodes where the ISRs for the said I/O devices reside.).

Kiick fails to teach where one or more of the nodes are memoryless.

Examiner takes Official Notice that nodes in a NUMA system may or may not
335 contain memory. Rowlands et al. evidences this (col. 4, lines 1 – 2).

Therefore, it would have been obvious to one of ordinary skill in this art at the time of invention by Applicant to provide memoryless nodes within the system of Kiick, since this is routine within NUMA systems.

Kiick fails to teach assigning the interrupts for the I/O devices to a first node
340 having a cache, memory, and at least one processor and assigning the interrupt for the I/O device to the first node. Similarly, Kiick fails to teach repeating the assignment for the second and third nodes, respectively, if the first node to which the device is connected does not have a cache, memory, and at least one processor.

Examiner takes Official Notice that it would be obvious to assign the interrupt to a
345 node which has a processor and memory, since these elements are required to service an interrupt (as evidenced by Agatsuma et al.; col. 4, line 64 - col. 5, line 9; col. 6, lines 15 - 22; Drott et al.; col. 21, lines 15 - 22; col. 22, lines 44 - 58; col. 23, lines 22 - 30; and Fischer et al.; col. 9, lines 35 - 51). Examiner notes that all modern processor architectures contain a "memory" since they contain registers which are operated upon
350 internally. Cache is old and notoriously well known in the art for providing improved performance in computing systems (as evidenced by Chi et al.; col. 2, lines 29 - 36; and Fischer et al.; col. 1, line 31 - col. 2, line 40).

Therefore, it would have been obvious to one of ordinary skill in this art at the time of invention by Applicant to assign the interrupt for the I/O devices of Kiick to nodes
355 which contain processors and memory, since both of these components are known in the prior art to be required in order to service an interrupt. It would have been obvious

to use choose nodes which have cache, since it is well-established in this art that caching provides greater performance.

360

With regard to claim 22, Kiick teaches the medium of claim 216, wherein the means is for assigning the interrupts among the plurality of nodes further based on whether the nodes have processors and memories (Kiick describes assigning ISRs to processors which have associated memories; paragraph 14).

365

With regard to claim 23, Kiick teaches the medium of claim 21, wherein the interrupt-assignor is to assign, for each node, the interrupts for the devices that are performance critical and that have been assigned to the node among the processors of the node in a round-robin manner (where the interrupt-assignor may be interpreted as a dynamic interrupt distributor; Examiner interprets all of the I/O devices of the invention of Kiick to be performance critical, thereby necessitating the use of his invention to improve performance; Paragraphs 25, 26, 28).

370

375

With regard to claim 24, Kiick teaches the medium of claim 23, wherein the interrupt-assignor is to dynamically modify assignments of the interrupts among the nodes based on actual performance characteristics of the assignments, and, for each node having processors, to dynamically modify assignments of the interrupts that are performance critical and that have been assigned to the node among the processors of

the node based on actual performance characteristics of the assignments (where an
380 interrupt-assignor may be interpreted as a dynamic interrupt distributor; paragraphs 25,
28, 31).

With regard to claim 25, Kiick teaches a method comprising:

Assigning interrupts for a plurality of input/output (I/O) devices among a plurality
385 of nodes based on at least one factor selected from the set consisting of: the nodes to
which the I/O devices are connected; and the nodes at which interrupt service routines
for the I/O devices reside, where one or more of the nodes have processors and
memory (Paragraph 34 describes that interrupts should be assigned to the “closest”
processors, and not across node boundaries. Examiner interprets this to mean the
390 interrupts for the I/O devices should be assigned to nodes to which they are connected
or to nodes where the ISRs for the said I/O devices reside.);

For each node of the system, assigning the interrupts for the devices that are
performance critical and that have been assigned to the node among the processors of
the node in a round-robin manner (Examiner interprets all of the I/O devices of the
395 invention of Kiick to be performance critical, thereby necessitating the use of his
invention to improve performance; Paragraphs 25, 26, 28);

Dynamically modifying assignments of the interrupts among the nodes of the
system based on actual performance characteristics of the assignments (paragraphs
25, 28, 31);

400 For each node of the system, dynamically modifying assignments of the
interrupts that are performance critical and that have been assigned to the node among
the processors of the node based on actual performance characteristics of the
assignments (paragraphs 25, 28, 31).

Kiick fails to teach where one or more of the nodes are memoryless.

405 Examiner takes Official Notice that nodes in a NUMA system may or may not
contain memory. Rowlands et al. evidences this (col. 4, lines 1 – 2).

Therefore, it would have been obvious to one of ordinary skill in this art at the
time of invention by Applicant to provide memoryless nodes within the system of Kiick,
since this is routine within NUMA systems.

410 Kiick fails to teach assigning the interrupts for the I/O devices to a first node
having a cache, memory, and at least one processor and assigning the interrupt for the
I/O device to the first node. Similarly, Kiick fails to teach repeating the assignment for
the second and third nodes, respectively, if the first node to which the device is
connected does not have a cache, memory, and at least one processor.

415 Examiner takes Official Notice that it would be obvious to assign the interrupt to a
node which has a processor and memory, since these elements are required to service
an interrupt (as evidenced by Agatsuma et al.; col. 4, line 64 - col. 5, line 9; col. 6, lines
15 - 22; Drott et al.; col. 21, lines 15 - 22; col. 22, lines 44 - 58; col. 23, lines 22 – 30;
and Fischer et al.; col. 9, lines 35 - 51). Examiner notes that all modern processor
420 architectures contain a "memory" since they contain registers which are operated upon
internally. Cache is old and notoriously well known in the art for providing improved

performance in computing systems (as evidenced by Chi et al.; col. 2, lines 29 – 36; and Fischer et al.; col. 1, line 31 – col. 2, line 40).

Therefore, it would have been obvious to one of ordinary skill in this art at the
425 time of invention by Applicant to assign the interrupt for the I/O devices of Kiick to nodes
which contain processors and memory, since both of these components are known in
the prior art to be required in order to service an interrupt. It would have been obvious
to use choose nodes which have cache, since it is well-established in this art that
caching provides greater performance.

430

With regard to claim 26, Kiick teaches the method of claim 25, wherein assigning
the interrupts for the plurality of I/O devices among the plurality of nodes of the system
comprises, for each I/O device:

Where the node (Figure 1, items 102A, 102B) to which the I/O device (Figure 1,
435 items 110A, 110B) is connected has a cache (Paragraph 10), memory (Figure 1, items
108A, 108B), and at least one processor (Figure 1, items 106A, 106B), assigning the
interrupt for the I/O device to the node to which the I/O device is connected;

Otherwise, where the node at which the interrupt service routine for the I/O
device resides has memory and at least one processor, assigning the interrupt for the
440 I/O device to the node at which the interrupt service routine for the I/O device resides
(Paragraph 34 describes that interrupts should be assigned to the “closest” processors,
and not across node boundaries. Examiner interprets this to mean the interrupts for the

I/O devices should be assigned to nodes to which they are connected or to nodes where the ISRs for the said I/O devices reside.).

445

With regard to claim 27, Kiick teaches the method of claim 25, wherein assigning the interrupts for the plurality of I/O devices among the plurality of nodes of the NUMA system further comprises, for each I/O device, otherwise, assigning the interrupt for the I/O device to one of the nodes having memory and at least one processor (Paragraph 23 describes each node containing memory (Figure 1, items 108A, 108B), and at least one processor (Figure 1, items 106A, 106B); paragraph 26).

450

With regard to claim 28, Kiick teaches the method of claim 25, wherein dynamically modifying the assignments of the interrupts among the nodes of the system comprises:

455

Measuring responsiveness of the node in processing the interrupt (paragraphs 27 – 30);

460

Kiick fails to explicitly teach assigning the interrupt to the node at which the interrupt service routine for the I/O device resides; measuring responsiveness of the node at which the interrupt service routine for the I/O device resides in processing the interrupt; and where the responsiveness of the node to which the I/O device is connected is better than the responsiveness of the node at which the interrupt service routine for the I/O device resides, reassigning the interrupt to the node to which the I/O device is connected.

465 Kiick does, however, teach that the dynamic interrupt distributor should be aware
of the system architecture, and re-assign interrupts to the “closest” processors
(paragraph 34). Therefore, it would have been obvious to one of ordinary skill in this art
at the time of invention by Applicant to give preference in assigning interrupts to nodes
at which the ISR for the device resides, or at which the I/O device itself resides, and
470 then to re-assign interrupts accordingly to which configuration produced better
performance as described in paragraphs 28 – 30).

 With regard to claim 29, Kiick teaches the method of claim 25, wherein
dynamically modifying the assignments of the interrupts among the nodes of the system
475 comprises, for each assignment of an interrupt for an I/O device to a node, where the
node is that at which the interrupt service routine for the I/O device resides:

 Measuring responsiveness of the node in processing the interrupt (paragraphs
27 – 30);

 Kiick fails to explicitly teach assigning the interrupt to the node to which the I/O
480 device is connected; measuring responsiveness of the node to which the I/O device is
connected in processing the interrupt; and where the responsiveness of the node at
which the interrupt service routine for the I/O device is connected is better than the
responsiveness of the node to which the I/O device is connected, reassigning the
interrupt to the node at which the interrupt service routine for the I/O device resides.

485 Kiick does, however, teach that the dynamic interrupt distributor should be aware
of the system architecture, and re-assign interrupts to the “closest” processors

(paragraph 34). Therefore, it would have been obvious to one of ordinary skill in this art at the time of invention by Applicant to give preference in assigning interrupts to nodes at which the I/O device is connected, or at which the ISR for the I/O device itself resides, and then to re-assign interrupts accordingly to which configuration produced better performance as described in paragraphs 28 – 30).

With regard to claim 30, Kiick teaches the method of claim 25, wherein for each node of the system having memory, dynamically modifying the assignments of the interrupts that are performance critical and that have been assigned to the node among the processors of the node comprises:

Measuring the responsiveness of the processors of the node in processing the interrupts assigned thereto (paragraphs 27, 28, 35);

Where a differential between a best responsiveness and a worst responsiveness is greater than a threshold (paragraph 28; where a threshold may be interpreted as a “large enough difference”);

Reassigning at least one of the interrupts assigned to the processor having the worst responsiveness to the processor having the best responsiveness (paragraphs 27 – 30, 35).

Response to Arguments

510 Applicant's arguments filed 4/8/2008 have been fully considered but they are not persuasive.

 Regarding Applicant's argument that there is no definitive selection of which node should have the interrupt assigned to it, the Examiner notes that the features upon which applicant relies (i.e., a definitive selection of which node should have the interrupt
515 assigned to it) are not recited in the rejected claim(s). Although the claims are interpreted in light of the specification, limitations from the specification are not read into the claims. See *In re Van Geuns*, 988 F.2d 1181, 26 USPQ2d 1057 (Fed. Cir. 1993). There is no step of selecting a predetermined order which determines which node is assigned to the interrupt first, rather only that a first, second and third node exist.

520 Regarding Applicant's argument that Kiick does not disclose assigning interrupts and modifying interrupt assignments, the Examiner notes that Kiick assigns the servicing of interrupts by assigning an interrupt service routine to a particular processor. Thus, Kiick teaches assigning interrupts, as in paragraph 12, lines 2 - 5, paragraph 24, lines 10 - 13, and paragraph 32, lines 9 - 12. Note that Kiick even describes, as would
525 be known to one of ordinary skill, to "re-distribute the interrupts to the new processors" in the aforementioned lines of paragraph 32. Thus, Kiick meets the limitation as claimed, given its broadest reasonable interpretation.

 Therefore, the Examiner cannot allow the claims.

530

Conclusion

THIS ACTION IS MADE FINAL. Applicant is reminded of the extension of time policy as set forth in 37 CFR 1.136(a).

A shortened statutory period for reply to this final action is set to expire THREE
535 MONTHS from the mailing date of this action. In the event a first reply is filed within
TWO MONTHS of the mailing date of this final action and the advisory action is not
mailed until after the end of the THREE-MONTH shortened statutory period, then the
shortened statutory period will expire on the date the advisory action is mailed, and any
extension fee pursuant to 37 CFR 1.136(a) will be calculated from the mailing date of
540 the advisory action. In no event, however, will the statutory period for reply expire later
than SIX MONTHS from the mailing date of this final action.

Any inquiry concerning this communication or earlier communications from the
examiner should be directed to MATTHEW D. SPITTLE whose telephone number is
(571)272-2467. The examiner can normally be reached on Monday - Friday, 9 - 5:30.

545 If attempts to reach the examiner by telephone are unsuccessful, the examiner's
supervisor, Mark Rinehart can be reached on 571-272-3632. The fax phone number for
the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for
550 published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a
555 USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

560 /M. D. S./
Examiner, Art Unit 2111

/Paul R. Myers/
Primary Examiner, Art Unit 2111